Order Optimal Coded Caching-Aided Multicast under Zipf Demand Distributions

Mingyue Ji*, Antonia M. Tulino†, Jaime Llorca† and Giuseppe Caire*
*Department of Electrical Engineering, University of Southern California
Email: {mingyuej, caire}@usc.edu
†Alcatel Lucent, Bell labs, Holmdel, NJ
Email: {a.tulino, jaime.llorca}@alcatel-lucent.com

arXiv:1402.4576v1 [cs.IT] 19 Feb 2014

***Abstract*—Caching and multicasting are two key technologies for reducing traffic load in content delivery networks. While uncoded caching based delivery schemes can offer competitive performance under skewed popularity distributions, the use of coded transmission schemes of increased complexity has been recently shown to significantly improve multicast efficiency under flatter popularity distributions, by exploiting coded multicast opportunities created by simple cooperative caching policies. In this paper, we consider a caching network with one source, hosting m files, connected to n destinations, each with a storage capacity of M files, via a shared link. Given that file requests follow a Zipf popularity distribution, our objective is to characterize the minimum average number of transmissions to satisfy all user demands in the information theoretic sense. We present both achievable coded caching-aided multicast schemes and outer bounds for this network configuration and show that as $m, n \to \infty$, for any M, the achievable average number of transmissions and the outer bound meets up to a constant factor.**

***Index Terms*—Multicast, caching, network coding, index coding, content distribution.**

I. Introduction

Content distribution services, and specially on-demand video streaming, are one of the key drivers of the exponential traffic growth experienced in today's networks [1]. A key characteristic of this type of traffic is the time-shifted nature of user requests for the same content, which reduces the opportunities for *naive multicasting* and impacts overall network load. Due to the increasing cost and scarcity of bandwidth resources, a promising approach for reducing network load is the use of relatively cheap caching resources close to end users.

The efficiency and throughput of different caching networks has been studied in recent works [2]–[8]. In [2], [3], Llorca *et al.* presented a formulation for the general content distribution problem (CDP), where nodes in an arbitrary network are allowed to cache, forward, replicate, and code information in order to deliver arbitrary user demands with minimum overall network cost. The authors showed an equivalence between the CDP and the network coding problem over a so-called caching-demand augmented graph, which proved the polynomial solvability of the CDP under uniform demands, and the hardness of the CDP under arbitrary demands. They further showed, via simulation, that for the case of a single bottleneck network, while low complexity uncoded delivery schemes can offer competitive performance for skewed popularity distributions, under flatter popularity distributions, the use of index coding based transmission schemes that exploit coded multicasting opportunities created by simple randomized popularity-based caching policies achieve significant average load reductions, at the expense of increased computational complexity [2].

In [4], [5], Ji *et al.* considered the one-hop Device-to-Device (D2D) wireless caching network, where user devices with limited storage capacity are allowed to communicate between each other under a protocol channel model [9]. By careful design of the caching configuration and the use of either coded [5] or uncoded [4] transmission schemes, the throughput of the D2D caching network is shown to scale as $\Theta\left(\max\left\{\frac{M}{m}, \frac{1}{n}\right\}\right)$, where m is the total number of files, n is the total number of users and M is the per user storage capacity. These results, show that when $nM > m$ the throughput of the D2D caching network can grow linearly with the cache size M, proving the large potential of caching networks.

The peak load or worst case throughput of the single bottleneck network was studied in [6], [7]. Interestingly, a throughput of $\Theta\left(\max\left\{\frac{M}{m}, \frac{1}{n}\right\}\right)$ is also shown to be the fundamental worst case throughput of this network. However, since worst case demands happen rarely in practice, it is usually more relevant to look at average performance when content requests follow a popularity distribution. In this paper, we study the characterization of the average rate in the single bottleneck caching network where the demand is characterized by a Zipf popularity distribution [10]. In a simultaneous work presented in [8], the authors designed an achievable scheme for the expected rate minimization problem under arbitrary popularity distributions that, although shown to outperform traditional local caching policies for certain values of the system parameters, could not guarantee order optimality. In this work, we present an alternative achievable scheme that allows coding over the full set of requested objects, such that when the requests follow a Zipf distribution, order optimality can be guaranteed in the information theoretic sense independent of the system parameters. Moreover, via simulations, we demonstrate significant gains of the proposed coded caching-aided multicast scheme over best known uncoded delivery schemes and the scheme proposed in [8].

The paper is organized as follows. In Section II, we present the network model and the problem formulation. The achievable content distribution scheme for any arbitrary demand

is introduced in Section III. In Section IV, we prove the order optimality of the proposed scheme for Zipf demand distributions. Finally, we discuss results and conclude the paper in Section V. Due to space limitations, all proofs are omitted and can be found in [11].

II. Network Model and Problem Formulation

We consider a single bottleneck network in which a content source (e.g., data center server, base station, cache server) with access to a content library containing m files is connected through a shared link to n users (e.g., user terminals, access points, small cells, cache servers) each with a cache of size M files.[1] We assume each file is partitioned into B equal-size packets and the channel between the content source and all the users follows a shared deterministic model.

Let file requests follow the independent reference model (IRM) (*i.e.,* requests are independent and identically distributed across users and over time), and happen according to a popularity distribution $\mathbf{q} = [q_f]_{f=1}^m$. The goal is to design a content distribution scheme that determines the information to be cached by each user and the information to be transmitted over the shared link for every possible demand, such that all demands are satisfied and the expected rate[2] over the shared link is minimized.[3] We remark that under IRM, an optimal content distribution scheme does not require updating the caching configuration, and hence can be designed as a one-time offline configuration (pre-fetching) of the network caches and a transmission scheme for any possible demand.

Note that our problem is an instance of the coded content distribution problem (CDP) presented in [2] for the specific case of the single bottleneck network. As shown in [2], [3], the content distribution problem is equivalent to the network coding problem (NCP) in a caching-demand augmented graph, and hence NP-Hard for the general arbitrary demand setting. Moreover, note that for a given (uncoded) caching and demand configuration, finding the optimal transmission scheme in the single bottleneck CDP is equivalent to solving an index coding problem (ICP) [12] with side information given by the chosen caching configuration.

In the following section, we present and achievable scheme for the CDP in the single bottleneck network (also referred to as the caching problem) based on a (distributed) uncoded randomized popularity-based caching policy and a (centralized) index coding based transmission scheme, which is first order optimal in terms of the expected transmission rate, as shown in Section IV .

III. Achievable Scheme

A. Transmission scheme

As mentioned earlier, finding a transmission scheme for the CDP in the single bottleneck network is equivalent to finding and index code with side information given by the caching configuration. One of the simplest achievable index coding schemes is the chromatic number of the complement of the side information graph [12], also referred to as the conflict graph [13]. Although shown to be suboptimal and in some specific graphs with unbounded multiplicative gap [14], we will show that for the caching problem, given the possibility to also configure the side (or cached) information, the chromatic number is sufficient to achieve order optimality.

Let $\mathcal{M}$ and $\mathcal{Q}$ denote the packet level caching and demand configurations, respectively, where $\mathcal{M}_{u,f}$ denotes the packets of file f cached at node u, and $\mathcal{Q}_{u,f}$ the packets of file f requested by node u. The (undirected) conflict graph $\mathcal{H}_{\mathcal{M},\mathcal{Q}}$ is constructed as follows:

- Consider each packet requested by each user as a single node, which means that even if the same packet is requested by W different users, they are represented as W different nodes in $\mathcal{H}_{\mathcal{M},\mathcal{Q}}$.
- Create an edge between node o_1 and o_2 if 1) they do not represent the same packet, and 2) o_1 is not available in the cache of the user requesting o_2, or o_2 is not available in the cache of the user requesting o_1.

The achievable index coding scheme is given by the minimum vertex coloring of the conflict graph $\mathcal{H}_{\mathcal{M},\mathcal{Q}}$. The scheme simply transmits the XOR of the packets with the same color. Therefore, given $\mathcal{M}$ and $\mathcal{Q}$, the total number of transmissions in terms of packets is given by the minimum number of colors or chromatic number $\chi(\mathcal{H}_{\mathcal{M},\mathcal{Q}})$, and the normalized transmission rate is given by $\chi(\mathcal{H}_{\mathcal{M},\mathcal{Q}})/B$. Observe that, by design, our transmission scheme allows coding over the full set of requested packets.

B. Caching scheme

We propose an uncoded, distributed, and randomized caching policy. The randomized nature, not only provides distributed operation, but also allows creating cooperation between users that even if not directly connected can effectively exchange information via coded multicast transmissions. We use Algorithm 1 to let each user distributedly fill their caches according to caching distribution $\mathbf{p} = [p_f]_{f=1}^m$, with $\sum_{f=1}^m p_f = 1$ and $p_f \leq 1/M, \forall f$.[4] Note that while each user caches the same amount $p_f MB$ packets of file f, the randomized nature of the algorithm makes each user cache different packets of the same file, which is key to maximize the amount of distinct packets of the same file cached by the collective set of users, and create coded multicast opportunities to be exploited by the coded transmission scheme.

Given that the transmission scheme is based on the chromatic number $\chi(\mathcal{H}_{\mathcal{M},\mathcal{Q}})$, our goal is to find the caching distribution $\mathbf{p}$ that minimizes the expected rate $\mathbb{E}[\chi(\mathcal{H}_{\mathcal{M},\mathcal{Q}})/B]$, where the expectation is over the random demand and caching

[1] We define a file as the information unit requested by the client application, e.g., entire files in file download applications or video chunks in video streaming applications.

[2] The rate is defined as the number of transmissions in terms of files, and the expectation is over all possible demands.

[3] The optimality is captured in the information theoretic sense.

[4] In order to avoid caching duplicated packets and violating capacity constraints.

Algorithm 1 Distributed Random Caching Algorithm

Require: $\mathbf{p} = [p_f]_{f=1}^m$

1: **for all** $f \in \mathcal{F}$ **do**
2: Cache a random subset of $p_f MB$ distinct packets of file f
3: **end for**
4: **return** $\mathcal{M}$

configurations, $\mathcal{M}$ and $\mathcal{Q}$, as

$$\begin{aligned} \underset{\mathbf{p}}{\text{minimize}} \quad & \mathbb{E}[\chi(\mathcal{H}_{\mathcal{M},\mathcal{Q}})/B] \\ \text{subject to} \quad & \sum_{f=1}^{m} p_f = 1. \end{aligned} \tag{1}$$

Due to the difficulty of obtaining a close-form expression for the chromatic number, we first upper bound the expected rate $\mathbb{E}[\chi(\mathcal{H}_{\mathcal{M},\mathcal{Q}})/B]$ using the following theorem.

Theorem 1: For any given m, n, M, and $\mathbf{q}$, a content distribution scheme that uses caching policy in Algorithm 1 with caching distribution $\{\mathbf{p} : p_f \leq 1/M, \forall f\}$, and chromatic number based index coding transmission, is achievable, and its expected rate, defined as $\bar{R}(\mathbf{p}, \mathbf{q})$, for B large enough, satisfies

$$\begin{aligned} \bar{R}(\mathbf{p}, \mathbf{q}) &\triangleq \mathbb{E}[\chi(\mathcal{H}_{\mathcal{M},\mathcal{Q}})/B] \\ &\leq \bar{R}^{\text{ub}}(\mathbf{p}, \mathbf{q}) \triangleq \min\{\psi(\mathbf{p}, \mathbf{q}), \bar{m}\}. \end{aligned} \tag{2}$$

In (2), $\bar{m}$ is the average number of requested files, and

$$\psi(\mathbf{p}, \mathbf{q}) = \sum_{\ell=1}^{n} \binom{n}{\ell} \sum_{f=1}^{m} \rho_{f,\ell} (1 - p_f M)^{n-\ell+1} (p_f M)^{\ell-1}, \tag{3}$$

where $\rho_{f,l}$ denotes the probability that file f is the file with least memory assignment among the files requested by a subset of ℓ users. That is, $\rho_{f,l} \triangleq \mathbb{P}\left(f = \arg\min_{j \in \mathcal{F}^\ell} p_j\right)$ with $\mathcal{F}^\ell$ denoting the set of files requested by a subset of ℓ users. (see [11] for details). □

Our proposed caching distribution is then given by

$$\mathbf{p}^* = \underset{\mathbf{p}: p_f \leq 1/M, \sum_f p_f = 1}{\text{argmin}} \psi(\mathbf{p}, \mathbf{q}). \tag{4}$$

C. Content distribution scheme

The proposed content distribution scheme hence comprises a caching policy, as described by Algorithm 1 with caching distribution $\mathbf{p}^*$ given by (4), and an index coding transmission scheme, as described in Section III-A. Hence, using Theorem 1, the achievable expected rate of our proposed scheme is given by $\bar{R}(\mathbf{p}^*, \mathbf{q})$ and is upper bounded as:

$$\bar{R}(\mathbf{p}^*, \mathbf{q}) \leq \bar{R}^{\text{ub}}(\mathbf{p}^*, \mathbf{q}). \tag{5}$$

Note also that for any other caching distribution $\mathbf{p} \neq \mathbf{p}^*$,

$$\bar{R}^{\text{ub}}(\mathbf{p}^*, \mathbf{q}) \leq \bar{R}^{\text{ub}}(\mathbf{p}, \mathbf{q}), \quad \forall \mathbf{p} \neq \mathbf{p}^*. \tag{6}$$

IV. Order optimality

The focus of this section is to prove that the proposed content distribution scheme, as defined in Section III-C, is order optimal in the sense that the multiplicative gap between its expected rate and the optimal expected rate is finite.

Consequently, letting $\bar{R}^{\text{lb}}(\mathbf{q})$ denote a lower bound (converse) of the expected rate over any caching and delivery scheme, we say the our proposed scheme is order optimal if there exists a caching distribution $\tilde{\mathbf{p}}$ and a positive fine constant $c \geq 1$, such that

$$\lim_{n,m \to \infty} \frac{\bar{R}^{\text{ub}}(\tilde{\mathbf{p}}, \mathbf{q})}{\bar{R}^{\text{lb}}(\mathbf{q})} = c. \tag{7}$$

To this end, we choose $\tilde{\mathbf{p}}$ satisfying

$$\begin{aligned} \tilde{p}_f &= \frac{1}{\widetilde{m}}, \quad f \leq \widetilde{m} \\ \tilde{p}_f &= 0, \quad f \geq \widetilde{m} + 1 \end{aligned} \tag{8}$$

where $\widetilde{m}$ is a function of the system parameters, $\widetilde{m}(m, n, M, \mathbf{q})$.

Observe that

$$\bar{R}^{\text{lb}}(\mathbf{q}) \leq \bar{R}(\mathbf{p}^*, \mathbf{q}) \leq \bar{R}^{\text{ub}}(\mathbf{p}^*, \mathbf{q}) \leq \bar{R}^{\text{ub}}(\tilde{\mathbf{p}}, \mathbf{q}). \tag{9}$$

Due to the complexity of analyzing order optimality for any popularity distribution, we focus on the most relevant distribution, the so-called Zipf distribution [10].

Let the file popularity $\mathbf{q}$ follow a Zipf distribution with $q_f = \frac{f^{-\alpha}}{\sum_{i=1}^{m} i^{-\alpha}}$, $\forall f$. We notice that the behavior of Zipf distribution is fundamentally different for the two regions of the Zipf parameter $\alpha < 1$ and $\alpha > 1$. In the following, we consider the two cases separately[5].

A. When $\alpha < 1$

Theorem 2: When $\alpha < 1$, the proposed achievable scheme is order optimal and

$$\lim_{n,m \to \infty} \frac{\bar{R}^{\text{ub}}(\tilde{\mathbf{p}}, \mathbf{q})}{\bar{R}^{\text{lb}}(\mathbf{q})} = c, \tag{10}$$

with $\tilde{\mathbf{p}}$ given by (8), $\widetilde{m}(m, n, M, \mathbf{q}) = m$, and

$$c = \frac{12}{(1 - \varepsilon_1) \min\left\{1 - e^{-1}, \frac{1}{2}(1 - \varepsilon_2)(1 - \alpha)\right\}}, \tag{11}$$

where ε_1 and ε_2 are arbitrarily small positive constants. □

B. When $\alpha > 1$

For this region of the Zipf parameter, we consider three different cases depending on the relationship between the number of users n and the library size m. In particular, we consider the cases of $n = \omega(m^\alpha)$, $n = \Theta(m^\alpha)$, and $n = o(m^\alpha)$.

[5]Due to space limitations, we do not consider the case of $\alpha = 1$ in this paper. However, order optimality results can be obtained by using a similar approach to the one used for $\alpha > 1$, and can be found in [11].

1) When $n = \omega\left(m^{\alpha}\right)$:

Theorem 3: When $\alpha > 1$ and $n = \omega\left(m^{\alpha}\right)$, the proposed achievable scheme is order optimal and

$$\lim_{n,m\to\infty} \frac{\bar{R}^{\mathrm{ub}}(\tilde{\mathbf{p}}, \mathbf{q})}{\bar{R}^{\mathrm{lb}}(\mathbf{q})} = c, \tag{12}$$

with $\tilde{\mathbf{p}}$ given by (8), $\widetilde{m}(m, n, M, \mathbf{q}) = m$, and

$$c = \frac{12}{(1-\varepsilon_3)}, \tag{13}$$

where ε_3 is an arbitrarily small positive constant. □

2) When $n = \Theta\left(m^{\alpha}\right) = \rho m^{\alpha}$:

Theorem 4: When $\alpha > 1$ and $n = \Theta\left(m^{\alpha}\right) = \rho m^{\alpha}$, the proposed achievable scheme is order optimal and

$$\lim_{n,m\to\infty} \frac{\bar{R}^{\mathrm{ub}}(\tilde{\mathbf{p}}, \mathbf{q})}{\bar{R}^{\mathrm{lb}}(\mathbf{q})} = c, \tag{14}$$

where $\tilde{\mathbf{p}}$ is given by (8), $\widetilde{m}(m, n, M, \mathbf{q})$ admits the following expressions as a function of the system parameters:

- if $\rho \geq 1$, $\widetilde{m}(m, n, M, \mathbf{q}) = m$
- if $0 < \rho < 1$,
  - if $M < 1$, $\widetilde{m}(m, n, M, \mathbf{q}) = \rho^{\frac{1}{\alpha}} m$
  - if $1 \leq M < \frac{1}{\rho}$, $\widetilde{m}(m, n, M, \mathbf{q}) = \rho^{\frac{1}{\alpha}+\zeta} m^{\alpha\zeta+1}$
  - if $M \geq \frac{1}{\rho}$, $\widetilde{m}(m, n, M, \mathbf{q}) = m$

with $\zeta = \min\left\{\frac{\frac{1}{\alpha}\log M}{\log \rho m^{\alpha}}, \frac{\frac{1}{\alpha}\log\left(\frac{1}{\rho}\right)}{\log \rho m^{\alpha}}\right\}$, and c is given by:

$$\begin{aligned} c \;=\; \max\Bigg\{ & \frac{12}{\sigma_1}, \frac{2}{\frac{\sigma_2\delta_2\rho_1(\alpha-1)}{\alpha^2}\left(1-\frac{\frac{\sigma_2\delta_2\rho_1(\alpha-1)}{\alpha^2}}{1-\frac{\sigma_2\delta_2\rho_1(\alpha-1)}{\alpha^2}}\right)}, \\ & \frac{4\alpha^2}{\sigma_2\delta_2\rho_1(\alpha-1)\left(1-\frac{1}{\frac{\alpha^2}{\sigma_2\delta_2\rho_1(\alpha-1)}-1}\right)}, \\ & \frac{2}{\frac{\sigma_2\delta_2\rho_1(\alpha-1)}{\alpha^2}\left(1-\frac{\sigma_2\delta_2\rho_1(\alpha-1)}{\alpha^2}\right)}, \frac{2}{\sigma_2} \Bigg\}, \end{aligned} \tag{15}$$

where $0 < \sigma_1 < 1 - \exp\left(-\frac{\delta_1(\alpha-1)}{\alpha}\right)$ is a positive constant, and all other positive constants $\rho_1, \sigma_2, \delta_2 \in (0,1)$ are such that every term in (15) is positive. □

3) When $n = o\left(m^{\alpha}\right)$:

Theorem 5: When $\alpha > 1$ and $n = o\left(m^{\alpha}\right)$, the proposed achievable scheme is order optimal and

$$\lim_{n,m\to\infty} \frac{\bar{R}^{\mathrm{ub}}(\tilde{\mathbf{p}}, \mathbf{q})}{\bar{R}^{\mathrm{lb}}(\mathbf{q})} = c, \tag{16}$$

where $\tilde{\mathbf{p}}$ is given by (8), $\widetilde{m}(m, n, M, \mathbf{q})$ admits the following expression as function of the system parameters:

- if $M < 1$, $\widetilde{m}(m, n, M, \mathbf{q}) = n^{\frac{1}{\alpha}}$
- if $1 \leq M < \frac{m^{\alpha}}{n}$, $\widetilde{m}(m, n, M, \mathbf{q}) = n^{\frac{1}{\alpha}+\zeta}$
- if $M \geq \frac{m^{\alpha}}{n}$, $\widetilde{m}(m, n, M, \mathbf{q}) = m$

with $\zeta = \min\left\{\frac{\frac{1}{\alpha}\log M}{\log n}, \frac{\log m}{\log n} - \frac{1}{\alpha}\right\}$, and c is given by:

$$\begin{aligned} c = \max\Bigg\{ & \frac{2}{\frac{\sigma_3\delta_3\rho_2(\alpha-1)}{\alpha^2}}, \frac{2}{\sqrt{2}-1}, \\ & \frac{1}{\frac{\sigma_4\delta_4(\alpha-1)}{\alpha}\left(1-\frac{\sigma_4\delta_4(\alpha-1)}{\alpha}\right)}, \\ & \frac{\sigma_5\delta_5(\alpha-1)}{\alpha\left(1-\exp\left(-\frac{(1-\delta_5)^2(\alpha-1)}{2\alpha^2}\right)\right)}, \\ & \frac{2\alpha}{\sigma_5\delta_5(\alpha-1)\left(1-\exp\left(-\frac{(1-\delta_5)^2}{2\alpha\sigma_5\delta_5}\right)\right)} \\ & \cdot\frac{1}{\left(1-\exp\left(-\frac{(1-\delta_5)^2\delta_5(\alpha-1)}{2\sigma_5}\right)\right)} \\ & \frac{4}{\left(1-\exp\left(-\frac{(1-\delta_5)^2}{2\alpha\sigma_5\delta_5}\right)\right)} \\ & \cdot\frac{1}{\left(1-\exp\left(-\frac{(1-\delta_5)^2\delta_5(\alpha-1)}{2\sigma_5}\right)\right)} \Bigg\}, \end{aligned} \tag{17}$$

where all constants $\rho_2, \sigma_3, \delta_3, \sigma_4, \delta_4, \sigma_5, \delta_5 \in (0,1)$ are such that all terms in (17) are positive. □

V. Discussion

In Section IV, we have proved that the proposed content distribution scheme, which uses random caching with caching distribution $\mathbf{p}^*$ and chromatic number based index coding transmission is order optimal under Zipf demand distributions. Furthermore, from Theorems 2-5, we have shown that for achieving order optimality, it is sufficient to partition the library into maximum two subsets, and allocate the entire storage capacity evenly to the files in the first subset (containing the $\tilde{m}$ most popular files). The order optimal rate is then achieved by performing index coding over the requested packets from the first subset, and transmitting the requested files from the second subset entirely and uncoded.

Moreover, from Theorem 2, we have shown that for the particular case of $\alpha < 1$, one subset, and hence uniform random caching, is sufficient for order optimality. Hence, the scaling law of the optimal expected rate for $\alpha < 1$ is given by $\Theta\left(\min\{\frac{m}{M}, n\}\right)$, which coincides with the optimal worst case rate in both the single bottleneck and the D2D caching networks, as shown in [4], [6]. However, we note that significant constant gains can be achieved by optimizing the caching configuration according to the popularity distribution even in the region of $\alpha < 1$. For example:

Theorem 6: When $\alpha < 1$ and $n = \Theta(m) = \beta m$ for some constant β, then as $m, n \to \infty$, if we choose $\tilde{p}$ as given in (8) where $\tilde{m}(m, n, M, \mathbf{q})$ is given by

- $\tilde{m}(m, n, M, \mathbf{q}) = \beta^* m$ with $\beta^* = (\beta(1-\alpha)M)^{\frac{1}{\alpha}}$, if $M \leq \frac{1}{\beta(1-\alpha)}$,
- $\tilde{m}(m, n, M, \mathbf{q}) = m$, if $M > \frac{1}{\beta(1-\alpha)}$,

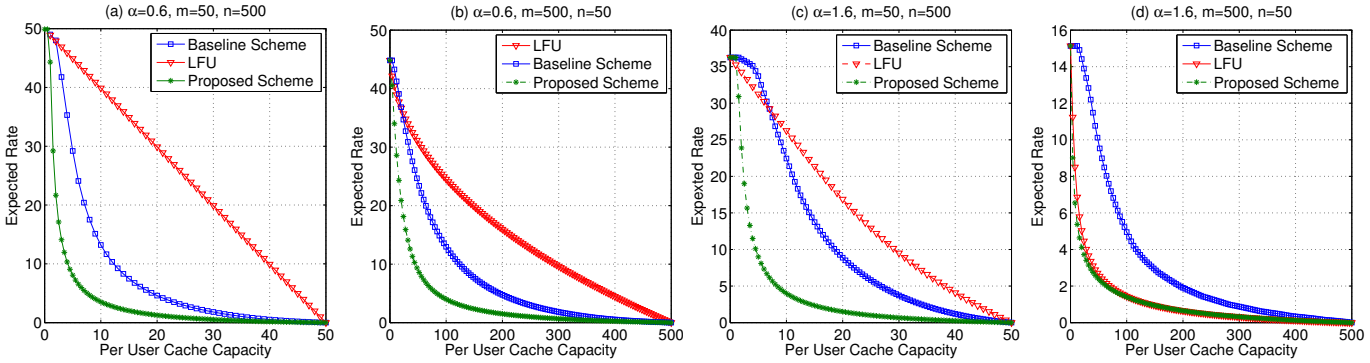

Fig. 1. The expected rate as a function of the per user cache capacity, M, for the proposed coded caching-aided multicast scheme, LFU, and the baseline scheme under the following values of the system parameters: $\alpha = \{0.6, 1.6\}$, $m = \{50, 500\}$ and $n = \{50, 500\}$.

then, the achievable expected rate satisfies:

$$\bar{R}^{\text{ub}}(\tilde{\mathbf{p}}, \mathbf{q}) \leq \min\left\{\left(\frac{\tilde{m}(m,n,M,\mathbf{q})}{M} - 1\right)\left(1 - e^{\left(-\frac{\beta}{\beta^* \alpha} M + o(1)\right)}\right) + (1 - {\beta^*}^{1-\alpha})\beta m\,,\ \left(\frac{m}{M} - 1\right)\left(1 - e^{-\beta M}\right) + o(m)\,,\ m\right\}.$$

□

From theorem 6, we can see that when n scales at the same speed as m, then once $M < \frac{1}{\beta(1-\alpha)}$, a constant factor improvement in the average rate can be achieved by exploiting the popularity distribution. Note that, although of harder analytical tractability, further constant gains can be achieved by using the optimized caching configuration $\mathbf{p}*$ in (4), as $\bar{R}^{\text{ub}}(\mathbf{p}^*, \mathbf{q}) \leq \bar{R}^{\text{ub}}(\tilde{\mathbf{p}}, \mathbf{q})$. Finally, from Theorem 4, observe that for $\alpha > 1$ and $n = \Theta(m^\alpha)$, the optimal expected rate can achieve an order gain with respect to the optimal peak rate.

Observe that unlike uncoded delivery schemes that transmit each non-cached packet separately, or the scheme proposed in [8], where files are grouped into subsets and coding is only performed within each subset, the key aspect of the order optimal schemes proposed in this paper, is the fact that coding is allowed within the entire set of requested packets. When treating different subsets of files separately, missed coding opportunities can significantly degrade multicast efficiency, as shown in Fig. 1, where we plot the upper bound of the rate achieved by our proposed scheme, $\bar{R}^{\text{ub}}(\tilde{\mathbf{p}}, \mathbf{q})$[6], compared to the rate achieved by the least frequently used (LFU)[7] caching algorithm and the scheme proposed in [8], which is referred to as the baseline scheme[8]. The expected rate is shown as a function of the per user cache capacity, M, for $n = \{50, 500\}$, $m = \{50, 500\}$, and $\alpha = \{0.6, 1.6\}$.

Observe that for all scenarios, the proposed scheme is able to significantly outperform both LFU and the baseline scheme. In particular, observe from Fig. 1c that for cache size 20% of the library size ($M = 10$), the proposed scheme achieves a factor improvement in expected rate of 5.6x with respect to the baseline scheme and 6.5x with respect to LFU. Note that LFU, while achieving close to optimal performance for high Zipf parameter and a library size greater than the number of users, as the number of users becomes larger than the library size and/or the Zipf parameter decreases, its performance significantly degrades. Also note how the baseline scheme can be very far from optimal, and even worse than LFU, for high Zipf parameter.

We remark that the improved performance and order optimality guarantees of the proposed schemes are based on the ability to 1) cache more packets of the more popular files, 2) maximize the amount of distinct packets of each file collectively cached, and 3) allow coded multicast transmissions within the full set of requested packets.

[6] $\bar{R}^{\text{ub}}(\tilde{\mathbf{p}}, \mathbf{q})$ is computed with $\tilde{\mathbf{p}}$ in (8) and $\tilde{m} = \operatorname{argmin} \bar{R}^{\text{ub}}(\tilde{\mathbf{p}}, \mathbf{q})$.

[7] LFU under IRM simply caches the M most popular files.

[8] The achievable rate for the scheme proposed in [8] is computed based on a grouping of the files to within a popularity factor of two, an optimization of the memory assigned to each group, and a separate coded transmission scheme for each group, as described in [8].

REFERENCES

[1] Cisco, "The Zettabyte Era-Trends and Analysis," 2013.

[2] J Llorca, A.M. Tulino, K Guan, J Esteban, M Varvello, N Choi, and D Kilper, "Network-coded caching-aided multicast for efficient content delivery," in *ICC, 2013 Proceedings IEEE*. IEEE, 2013.

[3] J Llorca and A.M. Tulino, "The content distribution problem and its complexity classification," Alcatel-Lucent technical report, 2013.

[4] M. Ji, G. Caire, and A. F. Molisch, "The throughput-outage tradeoff of wireless one-hop caching networks," *arXiv:1312.2637*, 2013.

[5] M. Ji, G. Caire, and A.F. Molisch, "Fundamental limits of distributed caching in d2d wireless networks," *arXiv:1304.5856*, 2013.

[6] M.A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *arXiv:1209.5807*, 2012.

[7] M. A. Maddah-Ali and U. Niesen, "Decentralized caching attains order-optimal memory-rate tradeoff," *arXiv:1301.5848*, 2013.

[8] U. Niesen and M.A. Maddah-Ali, "Coded caching with nonuniform demands," *arXiv:1308.0178*, 2013.

[9] P. Gupta and P.R. Kumar, "The capacity of wireless networks," *Information Theory, IEEE Transactions on*, vol. 46, no. 2, pp. 388–404, 2000.

[10] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and zipf-like distributions: Evidence and implications," in *INFOCOM'99. Proceedings. IEEE*. IEEE, 1999, vol. 1, pp. 126–134.

[11] M. Ji, A.M. Tulino, J. Llorca, and G. Caire, "Order optimal caching-aided multicast under zipf demand distributions," *In Preparation*.

[12] Y. Birk and T. Kol, "Informed-source coding-on-demand (iscod) over broadcast channels," 1998, IEEE.

[13] S. A. Jafar, "Topological interference management through index coding," *arXiv:1301.3106*, 2013.

[14] A. Blasiak, R. Kleinberg, and E. Lubetzky, "Index coding via linear programming," *arXiv:1004.1379*, 2010.